



TECHNICAL REPORT 3053  
November 2016

## Deep Learning Online Course

Katie Rainey

Approved for public release.

SSC Pacific  
San Diego, CA 92152-5001

**SSC Pacific**  
**San Diego, California 92152-5001**

---

**K. J. Rothenhaus, CAPT, USN**  
**Commanding Officer**

**C. A. Keeney**  
**Executive Director**

**ADMINISTRATIVE INFORMATION**

The work described in this report was performed by the Advanced Analysis Systems Branch (Code 56220) of the ISR Division (Code 56200), Space and Naval Warfare Systems Center Pacific (SSC Pacific), San Diego, CA. The Naval Innovative Science and Engineering (NISE) Program at SSC Pacific funded this project.

Released by  
E. Sharghi, Head  
Advanced Analysis Systems  
Branch

Under authority of  
C. A. Wilgenbusch, Head  
ISR Division

# **EXECUTIVE SUMMARY**

## **OBJECTIVE**

This technical report discusses the activities of the Naval Innovative Science and Engineering (NISE) project entitled Deep Learning Online Course, executed in fiscal year 2016 at Space and Naval Warfare Systems Center Pacific (SSC Pacific).

## **RESULTS**

The project was successful in training a large group of scientists and engineers in the topic of deep learning, a subfield of machine learning with an increasing number of application areas. Some difficulties were encountered throughout the course, which are documented herein.

## **RECOMMENDATIONS**

More courses of this sort should be held at SSC Pacific in the future. The lessons learned contained in this report should be taken into consideration when planning future courses.

# CONTENTS

<b>EXECUTIVE SUMMARY .....</b>	<b>iii</b>
<b>1. INTRODUCTION.....</b>	<b>1</b>
<b>2. COURSE STRUCTURE.....</b>	<b>2</b>
2.1 IN-PERSON MEETINGS .....	2
2.2 ONLINE COMMUNICATIONS.....	2
2.3 ADDITIONAL ACTIVITIES .....	2
2.4 MANAGEMENT .....	3
<b>3. IMPACT.....</b>	<b>4</b>
<b>4. DiSCUSSION.....</b>	<b>5</b>
4.1 BENEFITS.....	5
4.2 CHALLENGES .....	5
<b>5. FEEDBACK.....</b>	<b>7</b>
<b>REFERENCES.....</b>	<b>7</b>

## Figures

1. Number of people who participated in $n$ or more events .....	4
--	---

# 1. INTRODUCTION

CS231N is a course offered at Stanford University in Winter 2015 and again in Winter 2016 on the topic of “Convolutional Neural Networks for Visual Recognition.” The instructor was Prof. Fei-Fei Li, who is well known and is a leader in the computer vision community. All of the course materials were made available online [3], and the course notes and programming assignments are exceptionally well written and documented. Because of this, and because this is a particularly hot topic at the moment, the course became popular with many people outside of Stanford. There is a Reddit page for the course, and a computer science professor at the Naval Postgraduate School (NPS), Prof. Mathias Kolsch, used the materials to teach his own version of the course at NPS.

Meanwhile, topics relating to CS231N have become of greater and greater interest to many scientists and engineers at Space and Naval Warfare Systems Center Pacific (SSC Pacific). The machine learning community began organizing itself in 2012, which inspired a group of people to study an online course from Coursera.com on the basics of machine learning [4]. This group met approximately twice a month for seven months to discuss the course material and how it applied to projects at SSC Pacific. Between this course, the machine learning speaker series and workshop<sup>1</sup> [1], and a general increase in interest in machine learning, a large number of scientists and engineers had become interested in learning more about these related topics.

In FY2016, workforce development funding was provided by the Naval Innovative Science and Engineering (NISE) Program to support a group of SSC Pacific scientists and engineers who wished to take the CS231N course, led by principle investigator Katie Rainey (Code 56220). The funds covered primarily labor for the participants to study the material alongside their project work. This report documents the activities of the course along with some lessons learned.

---

<sup>1</sup>The machine learning series and workshop were funded in part under NISE Workforce Development in FY2014-15. The workshop was held in March 2015.

## 2. COURSE STRUCTURE

The Stanford course was given in 14 lectures, which were recorded on video in Winter 2016. In addition to the lecture slides and videos<sup>2</sup>, detailed course notes were made available online. There were also three homework assignments with starter Python code aimed at helping the student learn to code the algorithms taught in the course.

As there was no “expert” in this subject available to teach the course, the participants at SSC Pacific set out to study the material individually with support and motivation from the larger group. We held in-person meetings and also had several avenues for online collaboration.

### 2.1 IN-PERSON MEETINGS

The kickoff meeting was held on November 16, 2015. It was attended by 31 people, including several participants from the Pacific C4ISR Department (Code H) and SSC Atlantic dialing in.

Six “coding parties” were held, two for each programming assignment. The idea behind the parties was to get a bunch of people in a room (or online) at the same time to work on the assignments concurrently so that they could ask each other for help if necessary. Twenty-six unique participants attended the coding parties, which were all held in the Collaborative Innovation Lab (CoIL) in Building A33.

Upon finding that more people were focusing on the course notes rather than the programming assignments, we began a series of discussion groups where one person would work through some of the technical details from the course notes. Seven discussion groups were held with 25 unique attendees.

### 2.2 ONLINE COMMUNICATIONS

Regular communications with the course participants were made via an email list, and were always cross-posted to the project’s wiki space and to Fusion. The hashtag #deeplearning shared interesting links and information on Fusion. To share code, we used Spork, a Space and Naval Warfare Systems Command (SPAWAR) internal code repository, and similar repositories on the Internet such as Bitbucket and GitHub.

Additionally, a Slack<sup>3</sup> team was established to foster communications outside of SPAWAR-internal channels, which was useful for non-SPAWAR participants and those embedded at non-SPAWAR locations. Slack was mostly used to share links and to ask and answer questions about the course material and about deep learning in general. Professor Mathias Kolsch, who taught a course at Naval Postgraduate School (NPS) based on this material, joined the Slack team and answered technical questions that some participants asked. This interaction wouldn’t have been possible without some sort of external tool.

### 2.3 ADDITIONAL ACTIVITIES

The deep learning course inspired renewed activity among the Machine Learning Series. At least three<sup>4</sup> Machine Learning Series talks hosted in FY16 were organized in part with course funding. In April

---

<sup>2</sup>Several months after Stanford’s winter quarter completed, and mid-way through SPAWAR’s course, Stanford removed the course videos from [youtube.com](https://www.youtube.com) due to legal concerns. They remain hosted unofficially on [archive.org](https://archive.org).

<sup>3</sup>Slack is a commercial, free-to-use, cloud-based collaboration tool for teams. It’s essentially a glorified chat room.

<sup>4</sup>Other Machine Learning Series talks took place throughout the year, but they were not organized by the deep learning course organizers, and so their attendance metrics are not contained herein. Generally, if a technical talk takes place outside of a funded

Dean Lee (Code 54320) talked about his work doing predictive analytics on V-22 aircraft for the Comprehensive Automated Maintenance Environment Optimized/Readiness Integration Center (CAMEO/RIC) project. In May, Don Waagen from the Army's Aviation and Missile Research, Development, and Engineering Center (AMRDEC) shared thoughts on automatic target recognition algorithms in a talk entitled "Cargo Cults, Rabbit Holes, and the Most Important Questions." And in June, Vu Tran, a researcher with Booz Allen Hamilton, presented his work on applications of convolutional neural networks for image and video.

There was also an attempt this year to establish a "journal club" at which participants would take turns sharing the technical details of a paper that they found interesting or relevant to their work. Due to the time commitment required for the presenter, this idea has been slow to take off. Only two journal club meetings were held in FY16, but hopefully we can pick it back up next year. In March, Ben Migliori (Code 56150) discussed two papers on generative visualization and training of deep convolutional neural networks [6, 7] and how they apply to his own work. In May, Justin Mauger (Code 56150) shared work looking at imagery from the perspective of topology [2].

Funding was used to support attendance of five course participants—Katie Rainey (Code 56220), Alexander Corelli (Code 56220), Mark Bilinski (Code 56150), Chris Barngrover (Code 71710), and John Reeder (Code 53624)—at the 2016 IEEE Computer Vision and Pattern Recognition (CVPR) conference in Las Vegas, NV, in late June. Two other course participants—Michael Walton (Code 56150), Kris Gibson (Code 56430) and Josh Harguess (Code 56220)—also attended on other project funding. The attendees all gained knowledge to support their current and future projects, but in September we held a CVPR roundtable event to share what we learned with the rest of the deep learning community. With most of the CVPR attendees present, we shared some of the interesting research we saw presented and answered questions about the current state of computer vision research.

## **2.4 MANAGEMENT**

The course received \$98,500 in FY16 Naval Innovative Science and Engineering (NISE) Program funding. The bulk of that was spent on participant labor, with some also spent on conference fees and travel for CVPR. Monthly spending was tracked and reported to NISE Program Manager Robin Laird through monthly status reports. Course initiation documents were also provided, including a project management plan, an execution summary, a quad chart, and an environmental review. A mid-year review presentation was delivered in April 2016. Those participants who completed programming assignments shared their code on Spork or other code repositories. An end-of-year summary was delivered in September, and a poster was presented at a NISE poster session. All management documents, and links to all code repositories can be found on the course wiki page [5].

---

workforce development effort like the deep learning course or the Machine Learning Series several years back, no one will bother to take attendance.

### 3. IMPACT

To calculate participation metrics for the course, we consider 39 ways this course has provided to the SPAWAR community to engage with the deep learning community. These events include the course kickoff meeting, the seven discussion groups, the six coding parties, the CVPR roundtable, and six deep learning-related talks held this year. This also includes non-events such as whether someone charged to the project during a given month, whether they posted to Spork their code from one of the three course assignments, whether they follow the *#deeplearning* hashtag on Fusion, whether they signed up for the course mailing list, and whether they joined the course’s Slack team. This encompasses 66 total participants from SPAWAR, plus five Naval Research Enterprise Internship Program (NREIP) or Science and Engineering Apprentice Program (SEAP) interns, and several participants from NPS. Of course, some of these “events” indicate only very passive participation, but generally the more events a person participates in the more engaged in the course they can be considered to be. Figure 1 charts the number of people who participated in  $n$  or more events throughout the course. Thirty-two people participated in three or more events, which includes people who indicated real interest but perhaps did not have the time to engage fully in the course. Twenty people participated in seven or more events, which includes anyone who charged for seven months. Fifteen people participated in 13 or more events, which roughly captures the core active participants, and 10 extremely engaged people participated in 17 or more events. Overall, 19 people charged labor to the course, there were 70 total attendees (25 unique) to the discussion groups, and there were 59 total attendees (27 unique) to the coding parties. Seventy-five people total participated in the 38 events. Notably, 56 of those participants did so without funding.

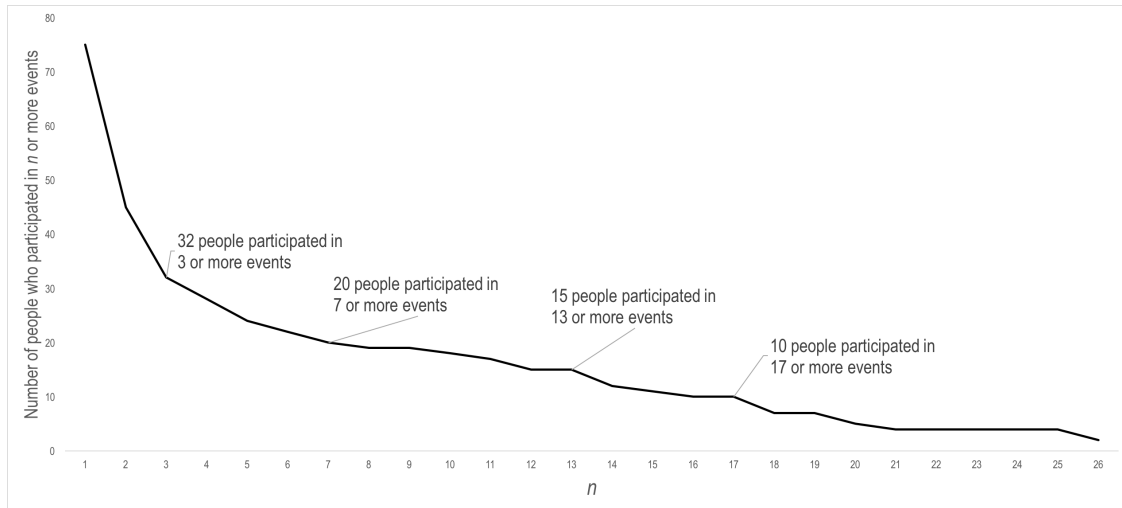


Figure 1. Number of people who participated in  $n$  or more events.



## 4. DISCUSSION

### 4.1 BENEFITS

Overall, there was strong, consistent participation in the course, and it gave good benefits to the S&T community for the cost. A couple dozen engineers and scientists now have a greater understanding of an important concept in machine learning, making them better equipped to design new systems for the warfighter, evaluate contractors' algorithms, and answer sponsors when they ask, "did you try deep learning on this problem?" This course has strengthened the S&T community and introduced scientists and engineers to people and projects outside of their respective divisions. It familiarized scientists and engineers with a technical area that is becoming relevant to more and more projects, and supported a culture of learning at the Center. It facilitated the hosting of technical talks of interest beyond just the course participants, and demonstrated to employees that the Center values their professional development. The events of the course sparked ideas and discussion about potential new projects, and provided a model for other similar courses in other technical areas.

### 4.2 CHALLENGES

**Course Material.** CS231n is an active undergraduate course at Stanford University, it is not an overview intended for a broad internet audience. It turns out that the course material is very technically challenging. Many participants found it difficult to keep pace with the course in only the small number of hours for which they received funding. The programming assignments were time consuming and especially hard for people without much experience coding in Python. Consequently, many people gave up on the programming and focused instead on understanding the course notes and watching the video lectures. This impacted attendance at coding parties and explains the lack of code committed to Spork.

**Coding Parties.** Attendance at the coding parties dropped after the first set, due to several factors. First is the fact that many people did not spend a lot of time working the programming assignments, as discussed above. Second is logistical issues that made coding parties not conducive to collaboration. CoIL is not properly set up for groups to code together. Not everyone worked on the course on a laptop, so they could not always bring their work with them to the parties. During some of the parties the computers in CoIL were not able to get on the internet, so the assignments couldn't be worked locally. Generally it was hard to work together when not everyone was working at the same pace. This became more true for the second and third assignments. Interesting discussions were had during the parties, but not much coding was accomplished.

**Long Distance Collaboration.** There was a lot of enthusiasm early on from participants not local to San Diego, including two students in Hawaii and a group at SSC Atlantic. But it was too logistically difficult to use DCS for the coding parties or discussion groups, and the Hawaii students in particular found it frustrating that they weren't able to fully participate. The Slack team was meant to bring the geographically distant together, and while it was used somewhat there simply wasn't enough activity to keep people engaged, with many people not signing up for the team in the first place.

**Funding.** Originally, funding was requested for 12 hours per month per participant over 7 months (84 hours total), but our low-ball offer was accepted instead which gave some participants only 6 hours (42 hours total). This was not enough. Participants found it difficult to carve only a small number of hours away from their regular project work. Spreading the course over seven months made it hard to dedicate the large chunks of time necessary to learn this difficult material. Because of this, despite not having enough

money, the project was perpetually underspent. Many participants only completed part of the course material in the time allotted, or did not gain as much detailed understanding as they could have.

It's not clear what the right solution is to helping the workforce learn new skills in our project-funded environment. Twenty hours per week for two months might allow greater dedication, but few people would be able to spare that kind of time away from their projects. Shortening the amount of time would make it hard for participants who had travel or other conflicts arise. Possibly, the solution is to organize smaller groups who could commit to a condensed schedule.

The idea behind our discussion groups was to ease the study burden for some participants. One volunteer would study the material from one lecture and present it to the group. This way people who had not found time to keep up with their reading could still attend and learn something. The discussion groups worked out relatively well. They were well attended and fostered lots of interesting dialog between people who had familiarity with using the concepts in practice and those who were completely unfamiliar with them. Unfortunately, we ran out of volunteers to lead the discussion groups.

## 5. FEEDBACK

This section contains feedback from several of the course participants.

Alan Li, Code H56F0:

Without something like this, I would have never tinkered with machine learning. I think this is a great way to get your feet wet in a foreign technical subject matter. The flexibility is a great safety net for those of us with other assignments that could have hard deadlines. This type of training is better than the typical multi-day/one-week training classes SPAWAR normally puts out because everyone can work at their own pace and take a deeper dive into the subject matter.

Lance Nakamoto, Code H56F0:

I would participate in more projects like this if the opportunity arose. Running through a college course is very good because we get lectures, structure, etc. The problem I encountered was being remote, I was not able to participate in the discussions/study group sessions. Slack ended up not being used much by people, so after a while it was just me and Alan (Zhao) in Hawaii, but we had conflicting schedules making it hard to collaborate.

Steve Hobbs, Code 52390:

The team learning experience was terrific, with many kinds of learning going on and good help from knowledgeable colleagues. I did not do everything that was in the course but I got out of it what I wanted most: (a) some programming experience for a few of the algorithms, (b) to learn something about what 'deep' learning is, and (c) to come to understand 'where we should be going from here' at SSC PAC.

## REFERENCES

1. Grady, D. 2013. “Mechanical Minds: Notes on Machine Learning.” <https://wiki.spawar.navy.mil/confluence/display/malsem/Mechanical+Minds+Home+Page>. Accessed 12 August 2016.
2. Lee, A. B., K. S. Pedersen, and D. Mumford. 2003. “The nonlinear statistics of high-contrast patches in natural images,” *International Journal of Computer Vision*, vol. 54, no. 1-3, pp. 83–103.
3. Li, F.-F. “CS231N: Convolutional Neural Networks for Visual Recognition.” <http://cs231n.stanford.edu/>. Accessed 15 September 2016.
4. Rainey, K. 2013. “Machine Learning MOOC Wiki Page.” <https://wiki.spawar.navy.mil/confluence/display/MOOC/Machine+Learning>. Accessed 12 August 2016.
5. Rainey, K. 2016. “Deep Learning Online Course Wiki Page.” <https://wiki.spawar.navy.mil/confluence/display/DLOC/CS231n%3A+Convolutional+Neural+Networks+for+Visual+Recognition>. Accessed 20 September 2016.
6. Zeiler, M. D. and R. Fergus. 2014. “Visualizing and understanding convolutional networks.” *Proceedings of the European Conference on Computer Vision* (pp. 818–833). Springer.
7. Zeiler, M. D., G. W. Taylor, and R. Fergus. 2011. “Adaptive deconvolutional networks for mid and high level feature learning.” *2011 International Conference on Computer Vision* (pp. 2018–2025). IEEE.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-01-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden to Department of Defense, Washington Headquarters Services Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p><b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</b></p>					
1. REPORT DATE (DD-MM-YYYY) November 2016		2. REPORT TYPE Final		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE  Deep Learning Online Course				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHORS  Katie Rainey				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SSC Pacific 53560 Hull Street San Diego, CA 92152-5001				8. PERFORMING ORGANIZATION REPORT NUMBER  TR 3053	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) SSC Pacific Naval Innovative Science and Engineering (NISE) Program SSC Pacific, 53560 Hull Street San Diego, CA 92152-5001				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release.					
13. SUPPLEMENTARY NOTES This is a work of the United States Government and therefore is not copyrighted. This work may be copied and disseminated without restriction.					
14. ABSTRACT  <p>This technical report discusses the activities of the Naval Innovative Science and Engineering (NISE) Program project entitled Deep Learning Online Course, executed in fiscal year 2016 at Space and Naval Warfare Systems Center Pacific (SSC Pacific).</p> <p>The project was successful in training a large group of scientists and engineers in the topic of deep learning, a subfield of machine learning with an increasing number of application areas. Some difficulties were encountered throughout the course, which are documented herein.</p> <p>More courses of this sort should be held at SSC Pacific in the future. The lessons learned contained in this report should be taken into consideration when planning future courses.</p>					
15. SUBJECT TERMS  Deep Learning Online Course; machine learning; workforce development					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Katie Rainey
U	U	U	U	15	19B. TELEPHONE NUMBER (Include area code) (619) 553-3472

## **INITIAL DISTRIBUTION**

84300	Library	(1)
56220	K.. Rainey	(1)

Defense Technical Information Center		
Fort Belvoir, VA 22060-6218		(1)

Approved for public release.



SSC Pacific  
San Diego, CA 92152-5001